7 Real Closed Fields and o-minimality

In this section, we will concentrate on the field of real numbers. Unlike algebraically closed fields, the theory of the real numbers does not have quantifier elimination in $\mathcal{L}_{\mathrm{r}} = \{+, 1, \cdot, 0, 1\}$, the language of rings. The proof of Corollary 6.5 shows that any field with quantifier elimination is strongly minimal, whereas in $\mathbb{R}$, if $\phi(x)$ is the formula $\exists z\ z^2 = x$, then ϕ defines an infinite coinfinite definable set. In fact, algebraically closed fields are the only infinite fields with quantifier eliination.

In fact, the ordering is the only obstruction to quantifier elimination. We will eventually analyze the real numbers in the language $\mathcal{L}_{\mathrm{or}} = \{+, -, \cdots, <, 0, 1\}$ and show that we have quantifier elimination in this language. Because the ordering $x < y$ is definable in the real field by the formula

$$\exists\ z\ (z \neq 0 \wedge x + z^2 = y),$$

any subset of $\mathbb{R}^n$ definable using an $\mathcal{L}_{\mathrm{or}}$-formula is already definable using an $\mathcal{L}_{\mathrm{r}}$-formula). We will see that quantifier elimination in $\mathcal{L}_{\mathrm{or}}$ leads us to a good geometric understanding of the definable sets.

We begin by reviewing some of the necessary algebraic background on ordered fields. All of the algebraic results stated in this chapter are due to Artin and Schreier.

Definition 7.1 We say that a field F is *orderable* if there is a linear order $<$ of F making $(F, <)$ an ordered field.

Although there are unique orderings of the fields $\mathbb{R}$ and $\mathbb{Q}$, orderable fields may have many possible orderings. The field of rational functions $\mathbb{Q}(X)$ has $2^{\aleph_0}$ distinct orderings. To see this, let x be any real number transcendental over $\mathbb{Q}$. The evaluation map $f(X) \mapsto f(x)$ is a field isomorphism between $\mathbb{Q}(X)$ and $\mathbb{Q}(x)$, the subfield of $\mathbb{R}$ generated by x. We can lift the ordering of the reals to an ordering $\mathbb{Q}(X)$ by $f(X) < g(X)$ if and only if $f(x) < g(x)$. Because $X < q$ if and only if $x < q$, choosing a different transcendental real would yield a different ordering. These are not the only orderings. We can also order $\mathbb{Q}(X)$ by making X infinite or infinitesimally close to a rational.

There is a purely algebraic characterization of the orderable fields.

Definition 7.2 We say that F is *formally real* if -1 is not a sum of squares.

In any ordered field all squares are nonnegative. Thus, every orderable field is formally real. The following result shows that the converse is also true.

Theorem 7.3 *If F is a formally real field, then F is orderable. Indeed, if $a \in F$ and $-a$ is not a sum of squares of elements of F, then there is an ordering of F where a is positive.*

Because the field of complex numbers is the only proper algebraic extension of the real field, the real numbers have no proper formally real algebraic extensions. Fields with this property will play a key role.

Definition 7.4 A field F is *real closed* if it is formally real with no proper formally real algebraic extensions.

Although it is not obvious at first that real closed fields form an elementary class, the next theorem allows us to axiomatize the real closed fields.

Theorem 7.5 *Let F be a formally real field. The following are equivalent.*

i) F is real closed.

ii) $F(i)$ is algebraically closed (where $i^2 = -1$).

iii) For any $a \in F$, either a or $-a$ is a square and every polynomial of odd degree has a root.

Corollary 7.6 *The class of real closed fields is an elementary class of $\mathcal{L}_{\mathrm{r}}$-structures.*

Proof We can axiomatize real closed fields by:

i) axioms for fields

ii) for each $n \geq 1$, the axiom

$$\forall x_1 \ldots \forall x_n \; x_1^2 + \ldots + x_n^2 + 1 \neq 0$$

iii) $\forall x \exists y \; (y^2 = x \vee y^2 + x = 0)$

iv) for each $n \geq 0$, the axiom

$$\forall x_0 \ldots \forall x_{2n} \exists y \; y^{2n+1} + \sum_{i=0}^{2n} x_i y^i = 0.$$

Although we can axiomatize real closed fields in the language of rings, we already noticed that we do not have quantifier elimination in this language. Instead, we will study real closed fields in $\mathcal{L}_{\mathrm{or}}$, the language of ordered rings. If F is a real closed field and $0 \neq a \in F$, then exactly one of a and $-a$ is a square. This allows us to order F by

$$x < y \text{ if and only if } y - x \text{ is a nonzero square.}$$

It is easy to check that this is an ordering and it is the only possible ordering of F.

Definition 7.7 We let RCF be the $\mathcal{L}_{\mathrm{or}}$-theory axiomatized by the axioms above for real closed fields and the axioms for ordered fields.

The models of RCF are exactly real closed fields with their canonical ordering. Because the ordering is defined by the $\mathcal{L}_{\mathrm{r}}$-formula

$$\exists z \; (z \neq 0 \wedge x + z^2 = y),$$

the next result tells us that using the ordering does not change the definable sets.

Proposition 7.8 *If F is a real closed field and $X \subseteq F^n$ is definable by an $\mathcal{L}_{\mathrm{or}}$-formula, then X is definable by an $\mathcal{L}_{\mathrm{r}}$-formula.*

Proof Replace all instances of $t_i < t_j$ by $\exists v\ (v \neq 0 \wedge v^2 + t_i = t_j)$, where t_i and t_j are terms occurring in the definition of X.

The next result suggests another possible axiomatization of RCF.

Theorem 7.9 *An ordered field F is real closed if and only if whenever $p(X) \in F[X]$, $a, b \in X$, $a < b$, and $p(a)p(b) < 0$, there is $c \in F$ such that $a < c < b$ and $p(c) = 0$.*

Definition 7.10 If F is a formally real field, a *real closure* of F is a real closed algebraic extension of F.

By Zorn's Lemma, every formally real field F has a maximal formally real algebraic extension. This maximal extension is a real closure of F.

The real closure of a formally real field may not be unique. Let $F = \mathbb{Q}(X)$, $F_0 = F(\sqrt{X})$, and $F_1 = F(\sqrt{-X})$. By Theorem 7.3, F_0 and F_1 are formally real. Let R_i be a real closure of F_i. There is no isomorphism between R_0 and R_1 fixing F because X is a square in R_0 but not in R_1. Thus, some work needs to be done to show that any ordered field $(F, <)$ has a real closure where the canonical order extends the ordering of F.

Lemma 7.11 *If $(F, <)$ is an ordered field, $0 < x \in F$, and x is not a square in F, then we can extend the ordering of F to $F(\sqrt{x})$.*

Proof We can extend the ordering to $F(\sqrt{x})$ by $0 < a + b\sqrt{x}$ if and only if
i) $b = 0$ and $a > 0$, or
ii) $b > 0$ and ($a > 0$ or $x > \frac{a^2}{b^2}$), or
iii) $b < 0$ and ($a < 0$ and $x < \frac{a^2}{b^2}$).

Corollary 7.12 *If $(F, <)$ is an ordered field, there is a real closure R of F such that the canonical ordering of R extends the ordering on F.*

Proof

By successive applications of Lemma 7.11, we can find an ordered field $(L, <)$ extending $(F, <)$ such that every positive element of F has a square root in L. We now apply Zorn's Lemma to find a maximal formally real algebraic extension R of L. Because every positive element of F is a square in R, the canonical ordering of R extends the ordering of F.

Although a formally real field may have nonisomorphic real closures, if $(F, <)$ is an ordered field there will be a unique real closure compatible with the ordering of F.

Theorem 7.13 *If $(F, <)$ is an ordered field, and R_1 and R_2 are real closures of F where the canonical ordering extends the ordering of F, then there is a unique field isomorphism $\phi : R_1 \to R_2$ that is the identity on F.*

Note that because the ordering of a real closed field is definable in $\mathcal{L}_{\mathrm{r}}$, ϕ also preserves the ordering. We often say that any ordered field $(F, <)$ has a unique real closure. By this we mean that there is a unique real closure that extends the given ordering.

Quantifier Elimination for Real Closed Fields

We are now ready to prove quantifier elimination.

Theorem 7.14 *The theory* RCF *admits elimination of quantifiers in $\mathcal{L}_{\mathrm{or}}$.*

Proof We use the quantifier elimination tests of §5. Suppose K and L are real closed ordered fields and A is a common substructure. Then A is an ordered integral domain. We extend the ordering on A to its fraction field to obtain an ordered subfield $F_0 \subseteq K \cap L$. Let F be the real closure of F_0. By uniqueness of real closures, F is isomorphic, as an ordered field, to the algebraic closure of F_0 inside K and L. Without loss of generality we may assume $F \subseteq K \cap L$.

It suffices then to show that if $\phi(v, \overline{w})$ is a quantifier-free formula, $\overline{a} \in F$, $b \in K$ and $K \models \phi(b, \overline{a})$, then there is $b' \in F$ such that $F \models \phi(b', \overline{a})$.

Note that

$$p(X) \neq 0 \leftrightarrow \ (p(\overline{X}) > 0 \vee -p(\overline{X}) > 0)$$

and

$$p(\overline{X}) \not> 0 \leftrightarrow (p(\overline{X}) = 0 \vee -p(\overline{X}) > 0).$$

With this in mind, we may assume that ϕ is a disjunction of conjunctions of formulas of the form $p(v, \overline{w}) = 0$ or $p(v, \overline{w}) > 0$. As in Theorem 6.1, we may assume that there are polynomials $p_1, \ldots, p_n$ and $q_1, \ldots, q_m \in F[X]$ such that

$$\phi(v, \overline{a}) \leftrightarrow \bigwedge_{i=1}^{n} p_i(v) = 0 \wedge \bigwedge_{i=1}^{m} q_i(v) > 0.$$

If any of the polynomials $p_i(X)$ is nonzero, then b is algebraic over F. Because F has no proper formally real algebraic extensions, in this case $b \in F$. Thus, we may assume that

$$\phi(v, \overline{a}) \leftrightarrow \bigwedge_{i=1}^{m} q_i(v) > 0.$$

The polynomial $q_i(X)$ can only change signs at zeros of q_i and all zeros of q_i are in F. Thus, we can find $c_i, d_i \in F$ such that $c_i < b < d_i$ and $q_i(x) > 0$ for all $x \in (c_i, d_i)$. Let $c = \max(c_1, \ldots, c_m)$ and $d = \min(d_1, \ldots, d_m)$. Then, $c < d$ and $\bigwedge_{i=1}^{m} q_i(x) > 0$ whenever $c < x < d$. Thus, we can find $b' \in F$ such that $F \models \phi(b', \overline{a})$.

Corollary 7.15 RCF *is complete, model complete and decidable. Thus* RCF *is the theory of* $(\mathbb{R}, +, \cdot, <)$ *and* RCF *is decidable.*

Proof By quantifier elimination, RCF is model complete.

Every real closed field has characteristic zero; thus, the rational numbers are embedded in every real closed field. Therefore, $\mathbb{R}_{\text{alg}}$, the field of real algebraic numbers (i.e., the real closure of the rational numbers) is a subfield of any real closed field. Thus, for any real closed field R, $\mathbb{R}_{\text{alg}} \prec R$, so $R \equiv \mathbb{R}_{\text{alg}}$.

In particular, $R \equiv \mathbb{R}_{\text{alg}} \equiv \mathbb{R}$.

Because RCF is complete and recursively axiomatized, it is decidable.

Semialgebraic Sets

Quantifier elimination for real closed fields has a geometric interpretation.

Definition 7.16 Let F be an ordered field. We say that $X \subseteq F^n$ is *semialgebraic* if it is a Boolean combination of sets of the form $\{\overline{x} : p(\overline{x}) > 0\}$, where $p(\overline{X}) \in F[X_1, \ldots, X_n]$.

By quantifier elimination, the semialgebraic sets are exactly the definable sets. The next corollary is a geometric restatement of quantifier elimination. It is analogous to Chevalley's Theorem (6.4) for algebraically closed fields.

Corollary 7.17 (Tarski–Seidenberg Theorem) *The semialgebraic sets are closed under projection.*

The next corollary is a typical application of quantifier elimination.

Corollary 7.18 *If $F \models RCF$ and $A \subseteq F^n$ is semialgebraic, then the closure (in the Euclidean topology) of A is semialgebraic.*

Proof We repeat the main idea of Lemma 1.26. Let d be the definable function

$$d(x_1, \ldots, x_n, y_1, \ldots, y_n) = z \text{ if and only if } z \geq 0 \wedge z^2 = \sum_{i=1}^{n} (x_i - y_i)^2.$$

The closure of A is

$$\{\overline{x} : \forall \epsilon > 0 \; \exists \overline{y} \in A \; d(\overline{x}, \overline{y}) < \epsilon\}.$$

Because this set is definable, it is semialgebraic.

We say that a function is semialgebraic if its graph is semialgebraic. The next result shows how we can use the completeness of RCF to transfer results from $\mathbb{R}$ to other real closed fields.

Corollary 7.19 *Let F be a real closed field. If $X \subseteq F^n$ is semialgebraic, closed and bounded, and f is a continuous semialgebraic function, then $f(X)$ is closed and bounded.*

Proof If $F = \mathbb{R}$, then X is closed and bounded if and only if X is compact. Because the continuous image of a compact set is compact, the continuous image of a closed and bounded set is closed and bounded.

In general, there are $\overline{a}, \overline{b} \in F$ and formulas ϕ and ψ such that $\phi(\overline{x}, \overline{a})$ defines X and $\psi(\overline{x}, y, \overline{b})$ defines $f(\overline{x}) = y$. There is a sentence Φ asserting:

> $\forall \overline{u}, \overline{w}$ [if $\psi(\overline{x}, y, \overline{w})$ defines a continuous function with domain $\phi(\overline{x}, \overline{u})$ and $\phi(\overline{x}, \overline{u})$ is a closed and bounded set, then the range of the function is closed and bounded].

By the remarks above, $\mathbb{R} \models \Phi$. Therefore, by the completeness of RCF, $F \models \Phi$ and the range of f is closed and bounded.

Model-completeness has several important applications. A typical application is Abraham Robinson's simple proof of Artin's positive solution to Hilbert's 17th problem.

Definition 7.20 Let F be a real closed field and $f(\overline{X}) \in F(X_1, \ldots, X_n)$ be a rational function. We say that f is *positive semidefinite* if $f(\overline{a}) \geq 0$ for all $\overline{a} \in F^n$.

Theorem 7.21 (Hilbert's 17th Problem) *If f is a positive semidefinite rational function over a real closed field F, then f is a sum of squares of rational functions.*

Proof Suppose that $f(X_1, \ldots, X_n)$ is a positive semidefinite rational function over F that is not a sum of squares. By Theorem 7.3, there is an ordering of $F(\overline{X})$ so that f is negative. Let R be the real closure of $F(\overline{X})$ extending this order. Then

$$R \models \exists \overline{v}\ f(\overline{v}) < 0$$

because $f(\overline{X}) < 0$ in R. By model-completeness

$$F \models \exists \overline{v}\ f(\overline{v}) < 0,$$

contradicting the fact that f is positive semidefinite.

We will show that quantifier elimination gives us a powerful tool for understanding the definable subsets of a real closed field.

Definition 7.22 Let $\mathcal{L} \supseteq \{<\}$. Let T be an $\mathcal{L}$-theory extending the theory of linear orders. We say that T is *o-minimal* if for all $\mathcal{M} \models T$ if $X \subseteq M$ is definable, then X is a finite union of points and intervals with endpoints in $M \cup \{\pm\infty\}$.

We can think of o-minimality as an analog of strong minimality for ordered structures. Strong minimality says that the only definable subsets in dimension one can be defined using only equality–i.e., the ones that can be defined in any structure. O-minimality says the only sets that can be defined in one dimension are the ones definable in any ordered structure.

Corollary 7.23 RCF*is an o-minimal theory.*

Proof Let $R \models$ RCF. We need to show that every definable subset of R is a finite union of points and intervals with endpoints in $R \cup \{\pm\infty\}$. By quantifier elimination, very definable subset of R is a finite Boolean combination of sets of the form

$$\{x \in R : p(x) = 0\}$$

and

$$\{x \in R : q(x) > 0\}.$$

Solution sets to nontrivial equations are finite, whereas sets of the second form are finite unions of intervals. Thus, any definable set is a finite union of points and intervals.

Next we will show that definable functions in one variable are piecewise continuous. The first step is to prove a lemma about $\mathbb{R}$ that we will transfer to all real closed fields.

Lemma 7.24 *If $f : \mathbb{R} \to \mathbb{R}$ is semialgebraic, then for any open interval $U \subseteq \mathbb{R}$ there is a point $x \in U$ such that f is continuous at x.*

Proof
case 1: There is an open set $V \subseteq U$ such that f has finite range on V.

Pick an element b in the range of f such that $\{x \in V : f(x) = b\}$ is infinite. By o-minimality, there is an open set $V_0 \subseteq V$ such that f is constantly b on V.

case 2: Otherwise.

We build a chain $U = V_0 \supset V_1 \supset V_2 \ldots$ of open subsets of U such that the closure $\overline{V}_{n+1}$ of V_{n+1} is contained in V_n. Given V_n, let X be the range of f on V_n. Because X is infinite, by o-minimality, X contains an interval (a, b) of length at most $\frac{1}{n}$. The set $Y = \{x \in V_n : f(x) \in (a, b)\}$ contains a suitable open interval V_{n+1}. Because $\mathbb{R}$ is locally compact,

$$\bigcap_{i=1}^{\infty} V_i = \bigcap_{i=1}^{\infty} \overline{V_i} \neq \emptyset.$$

If $x \in \bigcap_{i=1}^{\infty} V_i$, then f is continuous at x.

The proof above makes essential use of the completeness of the ordering of the reals. However, because the statement is first order, it is true for all real closed fields, by the completeness of RCF.

Corollary 7.25 *Let F be a real closed field and $f : F \to F$ is a semialgebraic function. Then, we can partition F into $I_1 \cup \ldots \cup I_m \cup X$, where X is finite and the I_j are pairwise disjoint open intervals with endpoints in $F \cup \{\pm\infty\}$ such that f is continuous on each I_j.*

Proof Let

$$D = \{x : F \models \exists \epsilon > 0\ \forall \delta > 0\ \exists y\ |x - y| < \delta \wedge |f(x) - f(y)| > \epsilon\}$$

be the set of points where f is discontinuous. Because D is definable, by o-minimality D is either finite or has a nonempty interior. By Corollary 7.23, D must be finite. Thus, $F \setminus D$ is a finite union of intervals on which F is continuous.

If F is real closed, then o-minimality tells us what the definable subsets of F look like. Definable subsets of F^n are also relatively simple.

Definition 7.26 We inductively define the collection of *cells* as follows.

- $X \subseteq F^n$ is a 0-cell if it is a single point.
- $X \subseteq F$ is a 1-cell if it is an interval (a, b), where $a \in F \cup \{-\infty\}$, $b \in F \cup \{+\infty\}$, and $a < b$.
- If $X \subseteq F^n$ is an n-cell and $f : X \to F$ is a continuous definable function, then $Y = \{(\overline{x}, f(\overline{x})) : \overline{x} \in X\}$ is an n-cell.
- Let $X \subseteq F^n$ be an n-cell. Suppose that f is either a continuous definable function from X to F or identically $-\infty$ and g is either a continuous definable function from X to F such that $f(\overline{x}) < g(\overline{x})$ for all $\overline{x} \in X$ or g is identically $+\infty$; then

$$Y = \{(\overline{x}, y) : \overline{x} \in X \wedge f(\overline{x}) < y < g(\overline{x})\}$$

is an $n + 1$-cell.

In a real closed field, every nonempty definable set is a finite disjoint union of cells. The proof relies on the following lemma.

Lemma 7.27 (Uniform Bounding) *Let $X \subseteq F^{n+1}$ be semialgebraic. There is a natural number N such that if $\overline{a} \in F^n$ and $X_{\overline{a}} = \{y : (\overline{a}, y) \in X\}$ is finite, then $|X_{\overline{a}}| < N$.*

Proof First, note that $X_{\overline{a}}$ is infinite if and only if there is an interval (c, d) such that $(c, d) \subseteq X_{\overline{a}}$. Thus $\{(\overline{a}, b) \in X : X_{\overline{a}} \text{ is finite}\}$ is definable. Without loss of generality, we may assume that for all $\overline{a} \in F^n$, $X_{\overline{a}}$ is finite. In particular, we may assume that

$$F \models \forall \overline{x} \forall c \forall d \neg [c < d \wedge \forall y (c < y < d \to y \in X_{\overline{a}})].$$

Consider the following set of sentences in the language of fields with constants added for each element of F and new constants $c_1, \ldots, c_n$. Let Γ be

$$\mathrm{RCF} + \mathrm{Diag}(F) + \left\{ \exists y_1, \ldots, y_m \ \left[\bigwedge_{i<j} y_i \neq y_j \wedge \bigwedge_{i=1}^{m} y_i \in X_{\overline{c}} \right] : m \in \omega \right\}$$

Suppose that Γ is satisfiable. Then, there is a real closed field $K \supseteq F$ and elements $\overline{c} \in K^n$ such that $X_{\overline{c}}$ is infinite. By model-completeness, $F \prec K$. Therefore

$$K \models \forall \overline{x} \forall c, d \ \neg [c < d \ \wedge \forall y \ (c < y < d \to y \in X_{\overline{a}})].$$

This contradicts the o-minimality of K. Thus, Γ is unsatisfiable and there is an N such that

$$\mathrm{RCF} + \mathrm{Diag}(F) \models \ \forall \overline{x} \ \neg \left(\exists y_1, \ldots, y_N \ \left[\bigwedge_{i<j} y_i \neq y_j \wedge \bigwedge_{i=1}^{N} y_i \in X_{\overline{x}} \right] \right).$$

In particular, for all $\overline{a} \in F^n, |X_{\overline{a}}| < N$.

We now state the Cell Decomposition Theorem and give the proof for subsets of F^2. In the exercises, we will outline the results needed for the general case.

Theorem 7.28 (Cell Decomposition) *Let $X \subseteq F^m$ be semialgebraic. There are finitely many pairwise disjoint cells $C_1, \ldots, C_n$ such that $X = C_1 \cup \ldots \cup C_n$.*

Proof (for $m = 2$) For each $a \in F$, let

$$C_a = \{x : \forall \epsilon > 0 \exists y, z \in (x - \epsilon, x + \epsilon)\ [(a, y) \in X \wedge (a, z) \notin X]\}.$$

We call C_a the *critical values* above a. By o-minimality, there are only finitely many critical values above a. By uniform bounding, there is a natural number N such that for all $a \in F$, $|C_a| \leq N$. We partition F into $A_0, A_1, \ldots, A_N$, where $A_n = \{a : |C_a| = n\}$.

For each $n \leq N$, we have a definable function $f_n : A_1 \cup \ldots \cup A_n \to F$ by $f_n(a) = n$th element of C_a. As above, $X_a = \{y : (a, y) \in X\}$.

For $n \leq N$ and $a \in A_n$, we define $P_a \in 2^{2n+1}$, the *pattern* of X above a, as follows.

If $n = 0$, then $P_a(0) = 1$ if and only if $X_a = F$. Suppose that $n > 0$.

$P_a(0) = 1$ if and only if $x \in X_a$ for all $x < f_1(a)$.

$P_a(2i - 1) = 1$ if and only if $f_i(a) \in X$.

For $i < n$, $P_a(2i) = 1$ if and only if $x \in X_a$ for all $x \in (f_i(a), f_{i+1}(a))$.

$P(2n) = 1$ if and only if $x \in X_a$ for all $x > f_n(a)$.

For each possible pattern $\sigma \in 2^{2n+1}$, let $A_{n,\sigma} = \{a \in A_n : P_a = \sigma\}$. Each $A_{n,\sigma}$ is semialgebraic. For each $A_{n,\sigma}$, we will give a decomposition of $\{(x, y) \in X : x \in A_{n,\sigma}\}$ into disjoint cells. Because the $A_{n,\sigma}$ partition F, this will suffice.

Fix one $A_{n,\sigma}$. By Corollary 7.25, we can partition $A_{n,\sigma} = C_1 \cup \ldots \cup C_l$, where each C_j is either an interval or a singleton and f_i is continuous on C_j for $i \leq n, j \leq l$. We can now give a decomposition of $\{(x, y) : x \in A_{n,\sigma}\}$ into cells such that each cell is either contained in X or disjoint from X.

For $j \leq l$, let $D_{j,0} = \{(x, y) : x \in C_j \text{ and } y < f(1)\}$.

For $j \leq l$ and $1 \leq i \leq n$, let $D_{j,2i-1} = \{(x, f_i(x)) : x \in C_j\}$.

For $j \leq l$ and $1 \leq i < n$, let $D_{j,2i} = \{(x, y) : x \in C_j, f_i(x) < y < f_{i+1}(x)\}$.

For $j \leq l$, let $D_{j,2n} = \{(x, y) : x \in C_j, y > f_n(x)\}$.

Clearly, each $D_{j,i}$ is a cell, $\bigcup D_{j,i} = \{(x, y) : x \in A_{n,\sigma}\}$, and each $D_{j,i}$ is either contained in X or disjoint from X. Thus, taking the $D_{j,i}$ that are contained in X, we get a partition of $\{(x, y) \in X : x \in A_{n,\sigma}\}$ into disjoint cells.

o-minimal Expansions of $\mathbb{R}$

The proofs above used very little about semialgebraic sets beyond o-minimality. Indeed, they would work in any o-minimal expansion of the real field. Indeed, there is a rich theory of definable sets in o-minimal expansions of the reals. We will survey some of the results in this section. For full details, see van den Dries book *Tame topology and o-minimal structures*.

Let $\mathcal{R} = (\mathbb{R}, +, \cdot, <, \ldots)$ be an o-minimal expansion of the reals, i.e., a structure obtained by adding extra structure to the reals such that $\mathrm{Th}(\mathcal{R})$ is o-minimal. Below by "definable" we will mean definable in $\mathcal{R}$.

Theorem 7.29 *Assume $\mathcal{R}$ is an o-minima expansion of $\mathbb{R}$.*

i) Every definable subset of $\mathbb{R}^n$ is a finite union of cells.

ii) If $f : X \to \mathbb{R}^n$ is definable, there is a finite partition of X into cells $X_1, \cup, X_n$ such that $f|X_i$ is continuous for each i. Indeed, for any $r \geq 0$, we can choose the partition such that $f|X_i$ is $\mathcal{C}^r$ for each i.

An easy consequence of ii) is that definable sets have only finitely many connected components. Much more is true, for example:

• Definable bounded sets can be definably triangulated.

• Suppose $X \subseteq \mathbb{R}^{n+m}$ is definable. For $a \in \mathbb{R}^m$ let

$$X_a = \{\overline{x} = (x_1, \ldots, x_n) \in \mathbb{R}^n : (\overline{x}, a) \in X\}.$$

There are only finitely many definable homeomorphism types for the sets X_a.

• (Curve selection) If $X \subseteq \mathbb{R}^n$ is definable and a is in the closure of X, then there is a continuous definable $f : (0, 1) \to X$ such that

$$\lim_{x \to 1} f(x) = a.$$

• If G is a definable group, then G is definably isomorphic to a Lie group.

• If we assume in addition that all definable functions are majorized by polynomials, then many of the metric properties of semialgebraic sets and asymptotic properties of semialgebraic functions also generalize.

Of course, this leads to the question: are there interesting o-minimal expansions of $\mathbb{R}$?

$\mathbb{R}_{\mathrm{an}}$ and subanalytic sets

Most of the results on o-minimal structures mentioned above were proved before we knew of any interesting o-minimal structures other than the real field. The first new example of an o-minimal theory was given by van den Dries.

Let $\mathcal{L}_{\mathrm{an}} = \mathcal{L} \cup \{\widehat{f} : \text{for some open } U \supset [0,1]^n,\ f : U \to \mathbb{R} \text{ is analytic}\}$.

We define $\widehat{f} : \mathbb{R}^n \to \mathbb{R}$ by

$$\widehat{f}(x) = \begin{cases} f(x) & x \in [0,1]^n \\ 0 & \text{otherwise.} \end{cases}$$

We let $\mathbb{R}_{\mathrm{an}}$ be the resulting $\mathcal{L}_{\mathrm{an}}$-structure. Denef and van den Dries proved that $\mathbb{R}_{\mathrm{an}}$ is o-minimal and that $\mathbb{R}_{\mathrm{an}}$ has quantifier elimination if we add a function

$$D(x, y) = \begin{cases} x/y & \text{if } 0 \leq |x| \leq |y| \\ 0 & \text{otherwise} \end{cases}$$

to the language. Quantifier elimination is proven by using the Weierstrass preparation theorem to replace arbitrary analytic functions of several variables by analytic functions that are polynomial in one of the variables. Tarski's elimination procedure is then used to eliminate this variable.

Denef and van den Dries also showed that if $f : \mathbb{R} \to \mathbb{R}$ is definable in $\mathbb{R}_{\mathrm{an}}$, then f is asymptotic to a rational function. In particular, although we can define the restriction of the exponential function to bounded intervals, we cannot define the exponential function globally. It is also impossible to define the sine function globally; for its zero set would violate o-minimality.

Although $\mathbb{R}_{\mathrm{an}}$ may seem unnatural, the definable sets form an interesting class.

We say that $X \subseteq \mathbb{R}^n$ is *semi-analytic* if for all x in $\mathbb{R}^n$ there is an open neighborhood U of x such that $X \cap U$ is a finite Boolean combination of sets $\{\overline{x} \in U : f(\overline{x}) = 0\}$ and $\{\overline{x} \in U : g(\overline{x}) > 0\}$ where $f, g : U \to \mathbb{R}$ are analytic. We say that $X \subseteq \mathbb{R}^n$ is *subanalytic* if for all x in $\mathbb{R}^n$ there is an open U and $Y \subset \mathbb{R}^{n+m}$ a bounded semianalytic set such that $X \cap U$ is the projection of Y into U. It is well known that subanalytic sets share many of the nice properties of semialgebraic sets.

If $X \subset \mathbb{R}^n$ is bounded, then X is definable in $\mathbb{R}_{\mathrm{an}}$ if and only if X is subanalytic. Indeed $Y \subseteq \mathbb{R}^n$ is definable in $\mathbb{R}_{\mathrm{an}}$ if and only if it is the image of a bounded subanalytic set under a semialgebraic map. Most of the known properties of subanalytic sets generalize to sets defined in any polynomial bounded o-minimal theory.

Exponentiation

The big breakthrough in the subject came in 1991. While quantifier elimination for $\mathbb{R}_{\mathrm{exp}}$ is impossible, Wilkie proved the next best thing.

Theorem 7.30 (Wilkie) *Let $\phi(x_1, \ldots, x_m)$ be an L_{exp} formula. Then there is $n \geq m$ and $f_1, \ldots, f_s \in \mathbb{Z}[x_1, \ldots, x_n, e^{x_1}, \ldots, e^{x_n}]$ such that $\phi(x_1, \ldots, x_n)$ is equivalent to*

$$\exists x_{m+1} \ldots \exists x_n \; f_1(x_1, \ldots, x_n, e^{x_1}, \ldots, e^{x_n}) = \ldots = f_s(x_1, \ldots, x_n, e^{x_1}, \ldots, e^{x_n}) = 0.$$

Thus every formula is equivalent to an existential formula (this property is equivalent to model completeness) and every definable set is the projection of an exponential variety.

Wilkie's proof depends heavily on the following special case of a theorem of Khovanski. Before Wilkie's theorem, Khovanski's result was the best evidence that $\mathbb{R}_{\mathrm{exp}}$ is o-minimal; indeed Khovanski's theorem is also the crucial tool needed to deduce o-minimality from model completeness.

Theorem 7.31 (Khovanski) *If $f_1, \ldots, f_m : \mathbb{R}^n \to \mathbb{R}$ are exponential polynomials, then $\{x \in \mathbb{R}^n : f_1(x) = \ldots f_n(x) = 0\}$ has finitely many connected components.*

If $X \subseteq \mathbb{R}$ is definable in $\mathbb{R}_{\text{exp}}$ then by Wilkie's Theorem there is an exponential variety $V \subseteq \mathbb{R}^n$ such that X is the projection of V. By Khovanski's Theorem V has finitely many connected components and X is a finite union of points and intervals. Thus $\mathbb{R}_{\text{exp}}$ is o-minimal.

Using the o-minimality of $\mathbb{R}_{\text{exp}}$ one can improve some of Khovanski's results on "fewnomials". From algebraic geometry we know that we can bound the number of connected components of a hypersurface in $\mathbb{R}^n$ uniformly in the degree of the defining polynomial. Khovanski showed that it is also possible to bound the number of connected component uniformly in the number of monomials in the defining polynomial. We will sketch the simplest case of this. Let $\mathcal{F}_{n,m}$ be the collection of polynomials in $\mathbb{R}[X_1, \dots, X_n]$ with at most m monomials. For $p \in \mathcal{F}_{n,m}$ let

$$V^+(p) = \{\overline{x} = (x_1, \dots, x_n) \in \mathbb{R}^n : \bigwedge_{i=1}^{n} x_i \geq 0 \wedge p(\overline{x}) = 0\}.$$

We claim that there are only finitely many homeomorphism types of $V^+(p)$ for $p \in \mathcal{F}_{n,m}$. Let $\Phi_{m,n}(x_1, \dots, x_n, r_{1,1} \dots, r_{1,n}, \dots, r_{m,1}, \dots, r_{m,n}, a_1, \dots, a_m)$ be the formula

$$\exists w_1 \dots, w_m \ ((\bigwedge_{i=1^m} e^{w_i} = x_i) \wedge \sum_{i=1}^{m} a_i \prod_{j=1}^{n} e^{w_i r_{i,j}} = 0).$$

We see that Φ expresses

$$\sum_{i=1}^{m} a_i \prod_{j=1}^{n} x_j^{r_{i,j}} = 0.$$

Let $X_{\overline{r},\overline{a}}$ denote the set of $\overline{x} \in \mathbb{R}^n$ such that $\Phi(\overline{x}, \overline{r}, \overline{a})$ holds. By o-minimality, $\{X_{\overline{r},\overline{a}} : \overline{r} \in \mathbb{R}^{mn}, \overline{a} \in \mathbb{R}^m\}$ represents only finitely many homeomorphism types.

In addition to answering the question of o-minimality, some headway has been made on the problem of decidability. Making heavy use of Wilkie's methods and Khovanski's theorem, Macintyre and Wilkie have shown that if Schanuel's Conjecture in is true then the first order theory of $\mathbb{R}_{\text{exp}}$ is decidable. Where Schanuel's Conjecture is the assertion that if $\lambda_1, \dots, \lambda_n$ are complex numbers linearly independent over $\mathbb{Q}$, then the transcendence degree of the field

$$\mathbb{Q}(\lambda_1, \dots, \lambda_n, e^{\lambda_1}, \dots, e^{\lambda_n})$$

is at least n.

Miller provided an interesting counterpoint to Wilkie's theorem. Using ideas of Rosenlicht he showed that if $\mathcal{R}$ is any o-minimal expansion of the real field that contains a function that is not majorized by a polynomial, then exponentiation is definable in $\mathcal{R}$.

Let $\mathcal{L}_{\text{an,exp}}$ be $\mathcal{L}_{\text{an}} \cup \{e^x\}$ and let $\mathbb{R}_{\text{an,exp}}$ be the real numbers with both exponentiation and restricted analytic functions. Using the Denef-van den Dries quantifier elimination for $\mathbb{R}_{\text{an}}$ and a mixture of model-theoretic and valuation theoretic ideas, van den Dries, Macintyre, and I were able to show that $\mathbb{R}_{\text{an,exp}}$ has quantifier elimination if we add log to the language. Using quantifier elimination and Hardy field style arguments (but avoiding the geometric type of arguments used by Khovanski) we were able to show that $\mathbb{R}_{\text{an,exp}}$ is o-minimal.

Since the language $\mathcal{L}_{\text{an,exp}}$ has size $2^{\aleph_0}$, one would not expect to give a simple axiomatization of the first order theory of $\mathbb{R}_{\text{an,exp}}$. Ressayre noticed that the model-theoretic analysis of $\mathbb{R}_{\text{an,exp}}$ uses very little global information about exponentiation. This observation leads to a "relative" axiomatization. The theory $Th(\mathbb{R}_{\text{an,exp}})$ is axiomatized by the theory of $\mathbb{R}_{\text{an}}$ and axioms asserting that exponentiation is an increasing homomorphism from the additive group onto the multiplicative group of positive elements that majorizes every polynomial.

Using this axiomatization and quantifier elimination one can show that any definable function is piecewise given by a composition of polynomials, exp, log, and restricted analytic functions on $[0, 1]^n$. For example, the definable function $f(x) = e^{e^x} - e^{x^2} - 3x$ is eventually increasing and unbounded. Thus for some large enough $r \in \mathbb{R}$ there is a function $g : (r, +\infty) \to \mathbb{R}$ such that $f(g(x)) = x$ for $x > r$. The graph of g is the definable set $\{(x, y) : x > r \text{ and } e^{e^y} - e^{y^2} - 3y = x\}$. Thus g is a definable function and there is some way to express g explicitly as a composition of rational functions, exp, log, and restricted analytic functions. In most cases it is in no way clear how to get these explicit representations of an implicitly defined function. One important corollary is that every definable function is majorized by an iterated exponential.